

Multi-channel Opportunistic Access: A Case of Restless Bandits with Multiple Plays

Sahand Haji Ali Ahmad, Mingyan Liu

Abstract—This paper considers the following stochastic control problem that arises in opportunistic spectrum access: a system consists of n channels where the state (“good” or “bad”) of each channel evolves as independent and identically distributed Markov processes. A user can select exactly k channels to sense and access (based on the sensing result) in each time slot. A reward is obtained whenever the user senses and accesses a “good” channel. The objective is to design a channel selection policy that maximizes the expected discounted total reward accrued over a finite or infinite horizon. In our previous work we established the optimality of a greedy policy for the special case of $k = 1$ (i.e., single channel access) under the condition that the channel state transitions are positively correlated over time. In this paper we show under the same condition the greedy policy is optimal for the general case of $k \geq 1$; the methodology introduced here is thus more general. This problem may be viewed as a special case of the restless bandit problem, with *multiple plays*. We discuss connections between the current problem and existing literature on this class of problems.

I. INTRODUCTION

We consider the following stochastic control problem: there are n uncontrolled Markov chains, each an independent, identically-distributed, two-state discrete-time Markov process. The two states will be denoted as state 1 and state 0 and the transition probabilities are given by p_{ij} , $i, j = 0, 1$.

The system evolves in discrete time. In each time instance, a user selects exactly k out of the n processes and is allowed to observe their states. For each selected process that happens to be in state 1 the user gets a reward; there is no penalty for selecting a channel that turns out to be state 0 but each such occurrence represents a lost opportunity because the user is limited to selecting only k of them. The ones that the user does not select do not reveal their true states. Our objective is to derive a selection strategy whose total expected discounted rewarded over a finite or infinite horizon is maximized.

This is a Markov decision process (or MDP) problem [?]. Furthermore, it is a partially observed MDP (or POMDP) problem [?] due to the fact that the states of the underlying Markov processes are not fully observed at all times and that as a consequence the system *state* as perceived by the user is in the form of a probability distribution, also commonly referred to as the *information state* of the system [?]. This problem is also an instance of the restless bandit problem with multiple plays [?], [?], [?]. More discussion on this literature is provided in section ??.

The application of the above problem abstraction to multichannel opportunistic access is as follows. Each Markov process represents a wireless channel, whose state transitions reflect dynamic changes in channel conditions caused by fading, interference, and so on. Specifically, we will consider state 1 as the “good” state, in which a user (or transmitter) can successfully communicate with a receiver; state 0 is the “bad” state, in which communication will fail. The channel state is assumed to remain constant within a single discrete time step. A multichannel system consists of n distinct channels. A user who wishes to use a particular channel at the beginning of a time step must first sense or probe the state of the channel, and can only transmit in a channel probed to be in the “good” state in the same time step. The user cannot sense and access more than k channels at a time due to hardware limitations. If all k selected channels turn out to be in the “bad” state, the user has to wait till the beginning of the next time step to repeat the selection process.

This model captures some of the essential features of multichannel opportunistic access as outlined above. On the other hand, it has the following limitations: the simplicity of the iid two-state channel model; the implicit assumption that channel sensing is perfect and the lack of penalty if the user transmits in a bad channel due to imperfect sensing; and the assumption that the user can select an arbitrary set of k channels out of n (e.g., it may only be able to access a contiguous block of channels due to physical layer limitations). Nevertheless this model does allow us to obtain analytical insights into the problem, and more importantly, some insight into the more general problem of restless bandits with multiple plays.

This model has been used and studied quite extensively in the past few years, mostly within the context of opportunistic spectrum access and cognitive radio networks, see for example [?], [?], [?], [?]. [?] studied the same problem and proved the optimality of the greedy policy in the special case of $k = 1, n = 2$, [?] proved the optimality of the greedy policy in the case of $k = n - 1$, while [?], [?] looked for provably good approximation algorithms for a similar problem. Furthermore, the indexability (in the context of Whittle’s heuristic index and indexability definition [?]) of the underlying problem was studied in [?], [?].

Our previous work [?] established the optimality of the greedy policy for the special case of $k = 1$ for arbitrary n and under the condition $p_{11} \geq p_{01}$, i.e., when a channel’s state transitions are positively correlated. In this sense, the results reported in the present paper is a direct generalization of results in [?], as we shall prove the optimality of the

This work is partially supported by CNS-0238035 and CCF-0910765. S. H. A. Ahmad and M. Liu are with the Dept. of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48105, {shajiali, mingyan}@eecs.umich.edu.

Multi-channel Opportunistic Access: A Case of Restless Bandits with Multiple Plays

Sahand Haji Ali Ahmad, Mingyan Liu

Abstract—This paper considers the following stochastic control problem that arises in opportunistic spectrum access: a system consists of n channels where the state (“good” or “bad”) of each channel evolves as independent and identically distributed Markov processes. A user can select exactly k channels to sense and access (based on the sensing result) in each time slot. A reward is obtained whenever the user senses and accesses a “good” channel. The objective is to design a channel selection policy that maximizes the expected discounted total reward accrued over a finite or infinite horizon. In our previous work we established the optimality of a greedy policy for the special case of $k = 1$ (i.e., single channel access) under the condition that the channel state transitions are positively correlated over time. In this paper we show under the same condition the greedy policy is optimal for the general case of $k \geq 1$; the methodology introduced here is thus more general. This problem may be viewed as a special case of the restless bandit problem, with *multiple plays*. We discuss connections between the current problem and existing literature on this class of problems.

I. INTRODUCTION

We consider the following stochastic control problem: there are n uncontrolled Markov chains, each an independent, identically-distributed, two-state discrete-time Markov process. The two states will be denoted as state 1 and state 0 and the transition probabilities are given by p_{ij} , $i, j = 0, 1$.

The system evolves in discrete time. In each time instance, a user selects exactly k out of the n processes and is allowed to observe their states. For each selected process that happens to be in state 1 the user gets a reward; there is no penalty for selecting a channel that turns out to be state 0 but each such occurrence represents a lost opportunity because the user is limited to selecting only k of them. The ones that the user does not select do not reveal their true states. Our objective is to derive a selection strategy whose total expected discounted rewarded over a finite or infinite horizon is maximized.

This is a Markov decision process (or MDP) problem [?]. Furthermore, it is a partially observed MDP (or POMDP) problem [?] due to the fact that the states of the underlying Markov processes are not fully observed at all times and that as a consequence the system *state* as perceived by the user is in the form of a probability distribution, also commonly referred to as the *information state* of the system [?]. This problem is also an instance of the restless bandit problem with multiple plays [?], [?], [?]. More discussion on this literature is provided in section V.

The application of the above problem abstraction to multichannel opportunistic access is as follows. Each Markov process represents a wireless channel, whose state transitions reflect dynamic changes in channel conditions caused by fading, interference, and so on. Specifically, we will consider state 1 as the “good” state, in which a user (or transmitter) can successfully communicate with a receiver; state 0 is the “bad” state, in which communication will fail. The channel state is assumed to remain constant within a single discrete time step. A multichannel system consists of n distinct channels. A user who wishes to use a particular channel at the beginning of a time step must first sense or probe the state of the channel, and can only transmit in a channel probed to be in the “good” state in the same time step. The user cannot sense and access more than k channels at a time due to hardware limitations. If all k selected channels turn out to be in the “bad” state, the user has to wait till the beginning of the next time step to repeat the selection process.

This model captures some of the essential features of multichannel opportunistic access as outlined above. On the other hand, it has the following limitations: the simplicity of the iid two-state channel model; the implicit assumption that channel sensing is perfect and the lack of penalty if the user transmits in a bad channel due to imperfect sensing; and the assumption that the user can select an arbitrary set of k channels out of n (e.g., it may only be able to access a contiguous block of channels due to physical layer limitations). Nevertheless this model does allow us to obtain analytical insights into the problem, and more importantly, some insight into the more general problem of restless bandits with multiple plays.

This model has been used and studied quite extensively in the past few years, mostly within the context of opportunistic spectrum access and cognitive radio networks, see for example [?], [?], [?], [?]. [?] studied the same problem and proved the optimality of the greedy policy in the special case of $k = 1, n = 2$, [?] proved the optimality of the greedy policy in the case of $k = n - 1$, while [?], [?] looked for provably good approximation algorithms for a similar problem. Furthermore, the indexability (in the context of Whittle’s heuristic index and indexability definition [?]) of the underlying problem was studied in [?], [?].

Our previous work [?] established the optimality of the greedy policy for the special case of $k = 1$ for arbitrary n and under the condition $p_{11} \geq p_{01}$, i.e., when a channel’s state transitions are positively correlated. In this sense, the results reported in the present paper is a direct generalization of results in [?], as we shall prove the optimality of the

This work is partially supported by CNS-0238035 and CCF-0910765. S. H. A. Ahmad and M. Liu are with the Dept. of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48105, {shajiali, mingyan}@eecs.umich.edu.

greedy policy under the same condition but for any $n \geq k \geq 1$. The main thought process used to prove this more general result derives from that used in [?]. However, there were considerable technical difficulties we had to overcome to reach the conclusion.

In the remainder of this paper we first formulate the problem in Section II, present preliminaries in Section III, and then prove the optimality of the greedy policy in Section IV. We discuss our work within the context of restless bandit problems in Section V. Section VI concludes the paper.

II. PROBLEM FORMULATION

As outlined in the introduction, we consider a user trying to access the wireless spectrum pre-divided into n independent and statistically identical channels, each given by a two-state Markov chain. The collection of n channels is denoted by \mathcal{N} , each indexed by $i = 1, 2, \dots, n$.

The system operates in discrete time steps indexed by t , $t = 1, 2, \dots, T$, where T is the time horizon of interest. At time t^- , the channels go through state transitions, and at time t the user makes the channel selection decision. Specifically, at time t the user selects k of the n channels to sense, the set denoted by $a^k \subset \mathcal{N}$.

For channels sensed to be in the “good” state (state 1), the user transmits in those channels and collects one unit of reward for each such channel. If none is sensed good, the user does not transmit, collects no reward, and waits until $t+1$ to make another choice. This process repeats sequentially until the time horizon expires.

The underlying system (i.e., the n channels) is not fully observable to the user. Specifically, channels go through state transition at time t^- (or anytime between $(t-1, t)$), thus when the user makes the channel sensing decision at time t , it does not have the true state of any channel at time t . Furthermore, upon its action (at time t^+) only k channels reveal their true states. The user’s action space at time t is given by the finite set $a^k(t) \subset \mathcal{N}$, where $a^k(t) = \{i_1, \dots, i_K\}$.

We know (see e.g., [?], [?], [?]) that a sufficient statistic of such a system for optimal decision making, or the *information state* of the system [?], [?], is given by the conditional probabilities of the state each channel is in given all past actions and observations. Since each channel can be in one of two states, we denote this information state by $\bar{\omega}(t) = [\omega_1(t), \dots, \omega_n(t)] \in [0, 1]^n$, where $\omega_i(t)$ is the conditional probability that channel i is in state 1 at time t given all past states, actions and observations¹. Throughout the paper $\omega_i(t)$ will be referred to as the information state of channel i at time t , or simply the channel probability of i at time t .

Due to the Markovian nature of the channel model, the future information state is only a function of the current information state and the current action; i.e., it is independent of past history given the current information state and action.

It follows that the information state of the system evolves as follows. Given that the state at time t is $\bar{\omega}(t)$ and action $a^k(t)$ is taken, $\omega_i(t+1)$ for $i \in a^k(t)$ can take on two values: (1) p_{11} if the observation is that channel i is in a “good” state; this occurs with probability $\omega_i(t)$; (2) p_{01} if the observation is that channel i is in a “bad” state; this occurs with probability $1 - \omega_i$. For any other channel $j \notin a^k(t)$, with probability 1 the corresponding $\omega_j(t+1) = \tau(\omega_j(t))$ where the operator $\tau : [0, 1] \rightarrow [0, 1]$ is defined as

$$\tau(\omega) := \omega p_{11} + (1 - \omega) p_{01}, \quad 0 \leq \omega \leq 1. \quad (1)$$

The objective is to maximize its total discounted expected reward over a finite horizon given in the following problem (P) (extension to infinite horizon is discussed in Section V):

$$(P): \quad \max_{\pi} J_T^{\pi}(\bar{\omega}) = \max_{\pi} E^{\pi} \left[\sum_{t=1}^T \beta^{t-1} R_{\pi_t}(\bar{\omega}(t)) | \bar{\omega}(1) = \bar{\omega} \right]$$

where $0 \leq \beta \leq 1$ is the discount factor, and $R_{\pi_t}(\bar{\omega}(t))$ is the reward collected under state $\bar{\omega}(t)$ when channels in the set $a^k(t) = \pi_t(\bar{\omega}(t))$ are selected.

The maximization in (P) is over the class of deterministic Markov policies². An admissible policy π , given by the vector $\pi = [\pi_1, \pi_2, \dots, \pi_T]$, is such that π_t specifies a mapping from the current information state $\bar{\omega}(t)$ to a channel selection action $a^k(t) = \pi_t(\bar{\omega}(t)) \subset \{1, 2, \dots, n\}$. This is done without loss of optimality due to the Markovian nature of the underlying system, and due to known results on POMDPs [?, Chapter 6].

III. PRELIMINARIES

The dynamic programming (DP) representation of problem (P) is given as follows:

$$\begin{aligned} V_T(\bar{\omega}) &= \max_{a^k \in \mathcal{N}, |a^k|=k} E[R_{a^k}(\bar{\omega})] \\ V_t(\bar{\omega}) &= \max_{a^k \in \mathcal{N}, |a^k|=k} \left(\sum_{i \in a^k} \omega_i + \beta \cdot \right. \\ &\quad \left. \sum_{l_i \in \{0,1\}, i \in a^k} \left(\prod_{i \in a^k} \omega_i^{l_i} (1 - \omega_i)^{1-l_i} \right) \cdot \right. \\ &\quad \left. V_{t+1}(p_{01}, \dots, p_{01}, \tau(\omega_j), p_{11}, \dots, p_{11}) \right), \quad (2) \\ &\quad t = 1, 2, \dots, T-1. \end{aligned}$$

In the last term, the channel state probability vector consists of three parts: a sequence of p_{01} ’s that represent those channels sensed to be in state 0 at time t and the length of this sequence is the number of l_i ’s equaling zero; a sequence of values $\tau(\omega_j)$ for all $j \notin a^k$; and a sequence of p_{11} ’s that represent those channels sensed to be in state 1 at time t and the length of this sequence is the number of l_i ’s equaling one. Note that the future expected reward is calculated by summing over all possible realizations of the k selected channels.

¹Note that it is a standard way of turning a POMDP problem into a classic MDP problem by means of the information state, the main implication being that the state space is now uncountable.

²A Markov policy is a policy that derives its action only depending on the current (information) state, rather than the entire history of states, see e.g., [?].

The value function $V_t(\bar{\omega})$ represents the maximum expected future reward that can be accrued starting from time t when the information state is $\bar{\omega}$. In particular, we have $V_1(\bar{\omega}) = \max_{\pi} J_T^{\pi}(\bar{\omega})$, and an optimal deterministic Markov policy exists such that $a = \pi_t^*(\bar{\omega})$ achieves the maximum in (3) (see e.g., [?] (Chapter 4)).

For simplicity of representation, we introduce the following notations:

- $p_{01}[x]$: this is the vector $[p_{01}, p_{01}, \dots, p_{01}]$ of length x ;
- $p_{11}[x]$: this is the vector $[p_{11}, p_{11}, \dots, p_{11}]$ of length x .
- We will use the notation:

$$q(l_1, \dots, l_k) := \prod_{1 \leq i \leq k} (\omega_i^{l_i} (1 - \omega_i)^{1-l_i})$$

for $l_1, \dots, l_k \in \{0, 1\}$. That is, given a vector of 0s and 1s (total of k elements), $q()$ is the probability that a set of k channels are in states given by the vector.

With the above notation, Eqn (3) can be written as

$$\begin{aligned} V_t(\bar{\omega}) &= \max_{a^k \in \mathcal{N}, |a^k|=k} \left(\sum_{i \in a^k} \omega_i + \beta \cdot \right. \\ &\quad \left. \sum_{l_i \in \{0,1\}, i \in a^k} q(l_1, \dots, l_k) \cdot \right. \\ &\quad \left. V_{t+1}(p_{01}[k - \sum l_i], \dots, \tau(\omega_j), p_{11}[\sum l_i]) \right). \end{aligned}$$

Solving (P) using the above recursive equation can be computationally heavy, especially considering the fact that $\bar{\omega}$ is a vector of probabilities. It is thus common to consider suboptimal policies that are easier to compute and implement. One of the simplest such heuristics is a greedy policy where at each time step we take an action that maximizes the immediate one-step reward. Our focus is to examine the optimality properties of such a simple greedy policy.

For problem (P), the greedy policy under state $\bar{\omega} = [\omega_1, \omega_2, \dots, \omega_n]$ is given by

$$a^k(\bar{\omega}) = \arg \max_{a^k \in \mathcal{N}, |a^k|=k} \sum_{i \in a^k} \omega_i. \quad (3)$$

That is, the greedy policy seeks to maximize the reward as if there were only one step remaining in the horizon. In the next section we investigate the optimality of this policy. Specifically, we will show that it is optimal in the case of $p_{11} \geq p_{01}$. This extends the earlier result in [?] that showed this to be true for the special case of $k = 1$.

IV. OPTIMALITY OF THE GREEDY POLICY

In this section we show that the greedy policy is optimal when $p_{11} \geq p_{01}$. The main theorem of this section is as follows.

Theorem 1: The greedy policy is optimal for Problem (P) under the assumption that $p_{11} \geq p_{01}$. That is, for $t = 1, 2, \dots, T$, $k \leq n$, and $\forall \bar{\omega} = [\omega_1, \dots, \omega_n] \in [0, 1]^n$, we have

$$V_t^k(\bar{\omega}; z^k(\bar{\omega})) \geq V_t^k(\bar{\omega}; a^k), \quad \forall a^k \in \mathcal{N}, \quad (4)$$

where $z^k(\bar{\omega})$ is the subset whose elements (indices) correspond to the k largest values in $\bar{\omega}$, and $V_t^k(\bar{\omega}; a^k)$ the expected value of action a^k followed by behaving optimally.

Below we present a number of lemmas used in the proof of this theorem. The first lemma introduces a notation that allows us to express the expected future reward under the greedy policy.

Lemma 1: There exist T n -variable functions, denoted by $W_t^k(\bar{\omega})$, $t = 1, 2, \dots, T$, each of which is a polynomial of order 1^3 and can be represented recursively in the following form:

$$\begin{aligned} W_T^k(\bar{\omega}) &= \sum_{n-1+1 \leq i \leq n} \omega_i \\ W_t^k(\bar{\omega}) &= \sum_{n-1+1 \leq i \leq n} \omega_i + \beta \cdot \\ &\quad \sum_{l_n, l_{n-1}, \dots, l_{n+k-1} \in \{0,1\}} q(l_n, \dots, l_{n+k-1}) \cdot \\ &\quad W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_i), \dots, \tau(\omega_{n-k}), p_{11}[\sum l_i]). \end{aligned}$$

The proof is easily obtained using backward induction on t given the recursive equation and noting that the mapping $\tau()$ is linear. The detailed proof is thus omitted for brevity.

A few remarks are in order on this function $W_t^k(\bar{\omega})$.

- Firstly, when $\bar{\omega}$ is given by an ordered vector $[\omega_1, \omega_2, \dots, \omega_n]$ with $\omega_1 \leq \omega_2 \leq \dots \leq \omega_n$, $W_t^k(\bar{\omega})$ is the expected total discounted future reward (from t to T) by following the greedy policy.

This follows from how the greedy policy works in the special case of $p_{11} \geq p_{01}$. Note that in this case the conditional probability updating function $\tau(\omega)$ is a monotonically increasing function, i.e., $\tau(\omega_1) \geq \tau(\omega_2)$ for $\omega_1 \geq \omega_2$. Therefore the ordering of channel probabilities is preserved among those that are not observed.

If a channel has been observed to be in state “1” (respectively “0”), its probability at the next step becomes $p_{11} \geq \tau(\omega)$ (respectively $p_{01} \leq \tau(\omega)$) for any $\omega \in [0, 1]$. In other words, a channel observed to be in state “1” (respectively “0”) will have the highest (respectively lowest) possible probability among all channels.

Therefore if we take the initial information state $\bar{\omega}(1)$, order the channels according to their probabilities $\omega_i(1)$, and sense the highest k channels (top k of the ordered list) with ties broken randomly, then following the greedy policy means that in subsequent steps we will keep a channel in its current position if it was sensed to be in state 1 in the previous slot; otherwise, it was observed to be in state 0 and gets thrown to the bottom of the ordered list. The policy then selects the next top most (or rightmost) k channels on this new ordered list. This procedure is essentially the same as that given in the recursive expression of $W()$.

- Secondly, when $\bar{\omega}$ is not ordered, $W_t^k()$ reflects a policy that simply goes down the list of channels by the order fixed in $\bar{\omega}$, while each time tossing the ones

³Each function W_t is affine in each variable, when all other variables are held constant.

observed to be 0 to the end of the list and keeping those observed to be 1 at the top of the list.

- iii) Thirdly, the fact that W_t^K is a polynomial of order 1 and affine in each of its elements implies that

$$\begin{aligned} & W_t^K(\omega_1, \dots, \omega_{n-2}, y, x) \\ & - W_t^K(\omega_1, \dots, \omega_{n-2}, x, y) \\ = & (x - y)[W_t^K(\omega_1, \dots, \omega_{n-2}, 0, 1) - \\ & W_t^K(\omega_1, \dots, \omega_{n-2}, 1, 0)] . \end{aligned}$$

Similar results hold when we change the positions of x and y . To see this, consider the above as two functions of x and y , each having an x term, a y term, an xy term and a constant term. Since we are only swapping the positions of x and y in these two functions, the constant term remains the same, and so does the xy term. Thus the only difference is the x term and the y term, as given in the above equation. This linearity result is used later in our proof.

The next lemma establishes a sufficient condition for the optimality of the greedy policy.

Lemma 2: Consider Problem (P) under the assumption that $p_{11} \geq p_{01}$. To show that the greedy policy is optimal at time t given that it is optimal at $t + 1, t + 2, \dots, T$, it suffices to show that at time t we have

$$\begin{aligned} & W_t^K(\omega_1, \dots, \omega_j, x, y, \dots, \omega_n) \\ \leq & W_t^K(\omega_1, \dots, \omega_j, y, x, \dots, \omega_n), \end{aligned} \quad (5)$$

for all $x \geq y$ and all $0 \leq j \leq n - 2$, with $j = 0$ implying $W_t^K(x, y, \omega_3, \dots, \omega_n) \leq W_t^K(y, x, \omega_3, \dots, \omega_n)$.

Proof: Since the greedy policy is optimal from $t + 1$ on, it is sufficient to show that selecting the best k channels followed by the greedy policy is better than selecting any other set of k channels followed by the greedy policy. If channels are ordered $\omega_1 \leq \dots \leq \omega_i \leq \dots \leq \omega_n$ then the reward of the former is precisely given by $W_t^K(\omega_1, \dots, \omega_n)$. On the other hand, the reward of selecting an arbitrary set a^k of k channels followed by acting greedily can be expressed as $W_t^K(\bar{a}^k, a^k)$, where \bar{a}^k is the (increasingly) ordered set of channels not included in a^k . It remains to show that if Eqn (5) is true then we have $W_t^K(\bar{a}^k, a^k) \leq W_t^K(\omega_1, \dots, \omega_n)$. This is easily done since the ordered list (\bar{a}^k, a^k) may be converted to $\omega_1, \dots, \omega_n$ through a sequence of switchings between two neighboring elements that are not increasingly ordered. Each such switch invokes (5), thereby maintaining the “ \leq ” relationship. \square

Lemma 3: For $0 \leq \omega_1 \leq \omega_2 \leq \dots \leq \omega_n \leq 1$, we have the following two inequalities for all $t = 1, 2, \dots, T$:

$$\begin{aligned} (A) : & 1 + W_t^K(\omega_2, \dots, \omega_n, \omega_1) \geq W_t^K(\omega_1, \dots, \omega_n) \\ (B) : & W_t^K(\omega_1, \dots, \omega_j, y, x, \omega_{j+3}, \dots, \omega_n) \geq \\ & W_t^K(\omega_1, \dots, x, y, \omega_{j+3}, \dots, \omega_n), \end{aligned}$$

where $x \geq y$, $0 \leq j \leq n - 2$, and $j = 0$ implies $W_t^K(y, x, \omega_3, \dots, \omega_n) \geq W_t^K(x, y, \omega_3, \dots, \omega_n)$.

This lemma is the key to our main result and its proof, which uses a sample path argument, highly instructive. It is however also lengthy, and for this reason has been relegated to the Appendix.

With the above lemmas, Theorem 1 is easily proven:

Proof of Theorem 1: We prove by induction on T . When $t = T$, the greedy policy is obviously optimal. Suppose it is also optimal for all times $t + 1, t + 2, \dots, T$, under the assumption $p_{11} \geq p_{01}$. Then at time t , by Lemma 2, it suffices to show that $W_t^K(\omega_1, \dots, \omega_j, x, y, \dots, \omega_n) \leq W_t^K(\omega_1, \dots, \omega_j, y, x, \dots, \omega_n)$ for all $x \geq y$ and $0 \leq j \leq n - 2$. But this is proven in Lemma 3. \square

V. DISCUSSION

While the formulation (P) is a finite horizon problem, the same result applies to the infinite horizon discounted reward case using standard techniques as we have done in our previous work [?], [?].

In the case of infinite horizon, the problem studied in this paper is closely associated with the class of multi-armed bandit problems [?] and restless bandit problems [?]. This is a class of problems where n controlled Markov chains (also called machines or arms) are activated (or played) one at a time. A machine when activated generates a state dependent reward and moves to the next state according to a Markov rule. A machine not activated either stays frozen in its current state (a rested bandit) or moves to the next state according to a possibly different Markov rule (a restless bandit). The problem is to decide the sequence in which these machines are activated so as to maximize the expected (discounted or average) reward over an infinite horizon.

The multi-armed bandit problem was originally solved by Gittins (see [?]), who showed that there exists an *index* associated with each machine that is solely a function of that individual machine and its state, and that playing the machine currently with the highest index is optimal. This index has since been referred to as the *Gittins index*. The remarkable nature of this result lies in the fact that it decomposes the n -dimensional problem into n 1-dimensional problems, as an index is defined for a machine independent of others. The restless bandit problem on the other hand was proven much more complex, and is PSPACE-hard in general [?]. Relatively little is known about the structure of its optimal policy in general. In particular, the Gittins index policy is not in general optimal [?].

When multiple machines are activated simultaneously, the resulting problem is referred to as multi-armed bandits with *multiple plays*. Again optimal solutions to this class of problems are not known in general. A natural extension to the Gittins index policy in this case is to play the machines with the highest Gittins indices (this will be referred to as the *extended Gittins index policy* below). This is not in general optimal for multi-armed bandits with multiple plays and an infinite horizon discounted reward criterion, see e.g., [?], [?]. However, it may be optimal in some cases, see e.g., [?] for conditions on the reward function, and [?] for an

undiscounted case where the Gittins index is always achieved at time 1. Even less is known when the bandits are restless, though asymptotic results for restless bandits with multiple plays were provided in [?] and [?].

The problem studied in the present paper is an instance of the restless bandits with multiple plays (in the infinite horizon case). Therefore what we have shown in this paper is an instance of the restless bandits problem with multiple plays, for which the extended Gittins index policy is optimal.

VI. CONCLUSION

In this paper we studied a stochastic control problem that arose in opportunistic spectrum access. A user can sense and access k out of n channels at a time and must select judiciously in order to maximize its reward. We extend a previous result where a greedy policy was shown to be optimal in the special case of $k = 1$ under the condition that the channel state transitions are positively correlated over time. In this paper we showed that under the same condition the greedy policy is optimal for the general case of $k \geq 1$. This result also contributes to the understanding of the class of restless bandit problems with multiple plays.

APPENDIX

Proof of Lemma 3: We would like to show

$$\begin{aligned} (A): \quad & 1 + W_t^k(\omega_2, \dots, \omega_n, \omega_1) \geq W_t^k(\omega_1, \dots, \omega_n) \\ (B): \quad & W_t^k(\omega_1, \dots, \omega_j, y, x, \omega_{j+3}, \dots, \omega_n) \geq \\ & W_t^k(\omega_1, \dots, x, y, \omega_{j+3}, \dots, \omega_n), \end{aligned}$$

where $x \geq y$, $0 \leq j \leq n-2$, and $j = 0$ implies $W_t^k(y, x, \omega_3, \dots, \omega_n) \geq W_t^k(x, y, \omega_3, \dots, \omega_n)$.

The two inequalities (A) and (B) will be shown together using an induction on t . For $t = T$, part (A) is true because $LHS = 1 + \omega_1 + \sum_{i=n-k+2}^n \omega_i \geq \omega_{n-k+1} + \sum_{i=n-k+2}^n \omega_i = RHS$. Part (B) is obviously true for $t = T$ since $x \geq y$.

Suppose (A) and (B) are both true for $t+1, \dots, T$. Consider time t , and we will prove (A) first. Note that in the next step, channel 1 is selected by the action on the LHS of (A) but not by the RHS, while channel $n-k+1$ is selected by the RHS of (A) but not by the LHS. Other than this difference both sides select the same set of channels indexed $n-k+2, \dots, n$. We now consider four possible cases in terms of the realizations of channels 1 and $n-k+1$.

Case (A.1): channels 1 and $n-k+1$ have the state realizations “0” and “1”, respectively.

We will use a sample-path argument. Note that while these two channels are not both observed by either side, the realizations hold for the underlying sample path regardless. In particular, even though the LHS does not select channel $n-k+1$ and therefore does not get to actually observe the realization of “1”, the fact remains that channel $n-k+1$ is indeed in state 1 under this realization, and therefore its future expected reward must reflect this. It follows that under this realization channel $n-k+1$ will have probability p_{11} for the next time step even though we did not get to observe the state 1. The same is true for the RHS. This argument applies to the other three cases and is thus not repeated.

Conditioned on this realization, the LHS and RHS are evaluated as follows (denoted as $\{LHS|_{(0,1)}\}$ and $\{RHS|_{(0,1)}\}$, respectively):

$$\begin{aligned} & \{LHS|_{(0,1)}\} \\ &= 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\ & \quad \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\ & \quad W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_2), \dots, \\ & \quad \tau(\omega_{n-k+1}) = p_{11}, p_{11}[\sum l_i]) ; \end{aligned}$$

$$\begin{aligned} & \{RHS|_{(0,1)}\} \\ &= 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\ & \quad \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\ & \quad W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_1) = p_{00}, \\ & \quad \tau(\omega_2), \dots, \tau(\omega_{n-k}), p_{11}[\sum l_i + 1]) \\ &= \{LHS|_{(0,1)}\} \end{aligned}$$

Case (A.2): channels 1 and $n-k+1$ have the state realizations “1” and “1”, respectively.

$$\begin{aligned} & \{LHS|_{(1,1)}\} \\ &= 1 + 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\ & \quad \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\ & \quad W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_2), \dots, \\ & \quad \tau(\omega_{n-k+1}) = p_{11}, p_{11}[\sum l_i + 1]) ; \end{aligned}$$

$$\begin{aligned} & \{RHS|_{(1,1)}\} \\ &= 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\ & \quad \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\ & \quad W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_1) = p_{11}, \\ & \quad \tau(\omega_2), \dots, \tau(\omega_{n-k}), p_{11}[\sum l_i + 1]) \\ &\leq 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\ & \quad \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\ & \quad W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\ & \quad p_{11}, p_{11}[\sum l_i + 1]) \\ &= \{LHS|_{(1,1)}\} - 1 \leq \{LHS|_{(1,1)}\} \end{aligned}$$

where the first inequality is due to the induction hypothesis of (B).

Case (A.3): channels 1 and $n - 1 + 1$ have the state realizations “0” and “0”, respectively.

$$\begin{aligned}
& \{RHS|_{(0,0)}\} \\
= & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_1) = p_{01}, \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\
& p_{11}[\sum l_i]) ; \\
& \{LHS|_{(0,0)}\} \\
= & 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\
& \tau(\omega_{n-k+1}) = p_{01}, p_{11}[\sum l_i]) \\
\geq & 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\
& p_{11}[\sum l_i], p_{01}) \\
\geq & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& \left(1 + W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_2), \dots, \tau(\omega_{n-k}), \right. \\
& \left. p_{11}[\sum l_i], p_{01})\right) \\
\geq & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}, p_{01}[k - \sum l_i], \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\
& p_{11}[\sum l_i]) \\
= & \{RHS|_{(0,0)}\}
\end{aligned}$$

where the first inequality is due to the induction hypothesis of (B), the last inequality due to the induction hypothesis of (A). Also, the second inequality utilizes the total probability over the distribution $q(l_{n-k+2}, \dots, l_n)$ and the fact that $\beta \leq 1$.

Case (A.4): channels 1 and $n - 1 + 1$ have the state

realizations “1” and “0”, respectively.

$$\begin{aligned}
& \{RHS|_{(1,0)}\} \\
= & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_1) = p_{11}, \tau(\omega_2), \dots, \\
& \tau(\omega_{n-k}), p_{11}[\sum l_i]) \\
& \{LHS|_{(1,0)}\} \\
= & 1 + 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\
& \tau(\omega_{n-k+1}) = p_{01}, p_{11}[\sum l_i + 1]) \\
\geq & 1 + 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\
& p_{11}[\sum l_i + 1], p_{01}) \\
\geq & 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& \left(1 + W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_2), \dots, \right. \\
& \left. \tau(\omega_{n-k}), p_{11}[\sum l_i + 1], p_{01})\right) \\
\geq & 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\
& p_{11}[\sum l_i + 1]) \\
\geq & 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], p_{11}, \tau(\omega_2), \dots, \tau(\omega_{n-k}), \\
& p_{11}[\sum l_i]) \\
= & 1 + \{RHS|_{(1,0)}\} \geq \{RHS|_{(1,0)}\}
\end{aligned}$$

where the first and last inequalities are due to the induction hypothesis of (B), the third due to the induction hypothesis

of (A).

With these four cases, we conclude the induction step of proving (A). We next prove the induction step of (B). We consider three cases in terms of whether x and y are among the top k channels to be selected in the next step.

Case (B.1): both x and y belong to the top k positions on both sides. In this case there is no difference between the LHS and RHS along each sample path, since both channels will be selected and the result will be the same.

Case (B.2): neither x nor y is among the top k positions on either side. This implies that $j \leq n - k - 2$. We have:

$$\begin{aligned}
& \text{LHS} \\
= & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_1), \dots, \tau(\omega_j), \\
& \tau(y), \tau(x), \tau(\omega_{j+3}), \dots, p_{11}[\sum l_i]) ; \\
& \text{RHS} \\
= & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_1), \dots, \tau(\omega_j), \\
& \tau(x), \tau(y), \tau(\omega_{j+3}), \dots, p_{11}[\sum l_i]) \\
\geq & \text{LHS}
\end{aligned}$$

where the last inequality is due to the monotonicity of $\tau()$ and the induction hypothesis of (B).

Case (B.3): exactly one of the two belongs to the the top k channels on each side. This implies that $j = n - k - 1$. By the linearity of the function W_t^k we have the following:

$$\begin{aligned}
& W_t^k(\omega_1, \dots, \omega_{n-k-1}, y, x, \omega_{n-k+2}, \dots, \omega_n) \\
& - W_t^k(\omega_1, \dots, \omega_{n-k-1}, x, y, \omega_{n-k+2}, \dots, \omega_n) \\
= & (x - y)(W_t^k(\omega_1, \dots, \omega_{n-k-1}, 0, 1, \omega_{n-k+2}, \dots, \omega_n) - \\
& W_t^k(\omega_1, \dots, \omega_{n-k-1}, 1, 0, \omega_{n-k+2}, \dots, \omega_n)) \quad (6)
\end{aligned}$$

However, we have

$$\begin{aligned}
& W_t^k(\omega_1, \dots, \omega_{n-k-1}, 1, 0, \omega_{n-k+2}, \dots, \omega_n) \\
= & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i], \tau(\omega_1), \dots, \tau(\omega_{n-k-1}), \\
& p_{11}, p_{11}[\sum l_i]) \\
\leq & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& \left(1 + W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_1), \dots, \right. \\
& \left. \tau(\omega_{n-k-1}), p_{11}[\sum l_i + 1], p_{01})\right) \\
\leq & \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& \left(1 + W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_1), \dots, \right. \\
& \left. \tau(\omega_{n-k-1}), p_{01}, p_{11}[\sum l_i + 1])\right) \\
\leq & 1 + \sum_{n-k+2 \leq i \leq n} \omega_i + \beta \cdot \\
& \sum_{l_{n-k+2}, \dots, l_n \in \{0,1\}} q(l_{n-k+2}, \dots, l_n) \cdot \\
& W_{t+1}^k(p_{01}[k - \sum l_i - 1], \tau(\omega_1), \dots, \tau(\omega_{n-k-1}), \\
& p_{01}, p_{11}[\sum l_i + 1]) \\
= & W_t^k(\omega_1, \dots, \omega_{n-k-1}, 0, 1, \omega_{n-k+2}, \dots, \omega_n)
\end{aligned}$$

Since $x \geq y$, we have $LHS \geq RHS$ in Eqn (6). This concludes the induction step of (B).